

SIMULATION AT TWO DIFFERENT LEVELS OF COMPLEXITY

Field of the Invention

5 The present invention relates to improvements in simulation. It is particularly applicable to the simulation of creatures for use in two different complexities of simulation.

10

Background of the Invention

As will be indicated later, although the present invention has broader application, it arose in connection with the
15 simulation of creatures. This area of art will now be discussed.

One of the problems with studying life sciences is that changes in ecosystems and the evolution of life forms
20 takes place over long periods of time. Consequently, life processes have been simulated in a variety of formats, thus allowing the study of the interaction of creatures and the emerging population patterns.

25 An animat is an artificial creature which exhibits an active and purposeful behaviour in an environment which presents it with problems - obstacles, navigation, hazards, goals. Animats are often inspired by biological organisms, which are observed to solve complex problems in
30 real physical environments. Indeed, animats may be designed to be as close as possible to the original biological organisms, so as to model the potential behaviour and/or evolution of such organisms.

It will be understood that use of the word "animat" is intended to mean not only artificial creatures that are directly inspired by real biological creatures, but also
5 imaginary animals and other autonomous agents such as real or imagined robots or "cyborgs" (part animal, part machine).

In order to make such simulations as accurate as possible,
10 it is desirable to lavish as much computational resource on each artificial creature as is available. The method of simulating an animat depends on the available computational resources. If a significant amount of processing power is available, it is possible to perform a
15 high-fidelity simulation (Hi-Fi Sim). Such a high-fidelity simulation will include the use of computationally expensive techniques for behaviour selection, sensing and interaction with the environment and other animats, e.g. use of a form of 3D rigid-body
20 physics simulation to model the environment and the various parts of the body of the animat; use of ray-tracing computer graphics to synthesize visual sensory input; the simulation of the transmission, absorption, and reflection of sound waves, etc.

25

If only a relatively small amount of computational power is available, then a relatively low-fidelity simulation (Lo-Fi Sim) would be utilised to ensure that the simulation progresses at a reasonable rate. Such a Lo-Fi
30 Sim will use computationally less expensive techniques, such as only using a two dimensional (2D) model of the environment, and employing a simplified model of the body of the animat e.g. treating it as a point mass with

predetermined locomotion characteristics (speed, rate of turn etc).

It is desirable that any predetermined type of animat will
5 function in a consistent manner regardless of the level of simulation.

Statement of the Invention

10

In a first aspect, the present invention provides a method of simulating a creature for use in two different complexities of simulation, the method comprising utilising a model of the creature that comprises at least
15 two portions, a first portion which contains functions for use in both of said different complexities of simulation; and a second portion comprising two alternative versions, a first version for use in one of said different complexities of simulation, and a second version for use
20 in the other of said different complexities of simulation.

In a further aspect, the invention provides a method of simulating a process at two different levels of complexity, the method comprising utilising a model that
25 comprises at least two portions, a first portion which contains functions for use in both of said different complexities of simulation; and a second portion comprising two alternative versions, a first version for use in one of said different complexities of simulation,
30 and a second version for use in the other of said different complexities of simulation.

It is thus appreciated that the present invention has broader application than to the simulation of creatures.

5 Brief Description of Drawings

For a better understanding of the invention, and to show how embodiments of the same may be carried into effect, reference will now be made, by way of example, to the
10 accompanying diagrammatic drawings in which:

Figure 1 shows a block diagram illustrating the simulation components that are independent of, and the components that are specific to, the different complexities of simulation, according to a preferred embodiment of the
15 present invention;

Figure 2 illustrates the different modes and associated transition probabilities between the modes of a high level brain, according to the preferred embodiment of the present invention; and

20 Figures 3A and 3B illustrate the generic application of aspects of the invention to simulation.

Detailed Description of Embodiments of the Invention

25

Firstly, a preferred embodiment of the invention for use in the simulation of creatures will be described. After this, a further, generic, embodiment of the invention will be discussed, together with its application to simulation
30 problems.

In the preferred embodiment of the present invention, a single computational resource is utilised to simulate a

relatively large animat population (e.g. 100,000's) in a Lo-Fi Sim, with an associated Hi-Fi Sim acting as a slave program to the Lo-Fi Sim.

The nature of the computational resource will not be
5 discussed in detail here, as conventional computational resources can be employed, as will be well understood by the skilled person. Any computational system with appropriate processor, memory and (where required) means of receiving instructions and displaying or otherwise
10 outputting data can be provided. Simultaneous simulation of a large number of entities is a computationally intense task, so real time simulation, for example, is likely to require reasonably high computational performance. As will be understood, however, aspects of the present
15 invention allow for most effective use of whatever computational resource is available for conducting the simulation.

The Hi-Fi Sim is engaged by the Lo-Fi Sim when a more
20 detailed simulation of the activities of one or more animats is required e.g. when an animat from that population is selected for closer inspection by a user, or alternatively when an animat requires more detailed simulation. For example, this could occur when the animat
25 encounters another animat and/or undergoes a change in its environment (e.g. encounters water or food) that could lead to an observable change in the behaviour of the animat.

30 It is desirable that there should be a high correspondence between the detailed model of the Hi-Fi Sim and more abstract representations used in the Lo-Fi Sim. For instance, it would clearly be undesirable if the Lo-Fi sim

were to simulate a particular animat as being attracted to another animat, whilst in the same situation the Hi-Fi Sim were to simulate the particular animat as being repulsed by this further animat. As the Hi-Fi Sim would be
 5 utilised to model the detailed interaction of these two animats, then this would lead to a cyclic behaviour, with the Lo-Fi Sim invoking the Hi-Fi Sim when the animats are close enough to interact, the Hi-Fi Sim then indicating that the animats would repulse each other, the Lo-Fi Sim
 10 indicating the animats would attract each other and invoking the Hi-Fi Sim again etc.

To overcome such problems, in the preferred embodiment the brain of each animat is separated into two parts: a high-
 15 level brain (HLB) and a low-level brain (LLB).

The HLB is a behaviour selection mechanism, that is arranged to use external (sensory) events and situations, as well as internal (e.g. simulated bio-chemical and
 20 energetic reaction levels) events and situations to decide what behaviour is appropriate at any given time. The HLB of an animat is used in both the Hi-Fi Sim and the Lo-Fi Sim. The internal operation of the HLB is identical in both simulations.

25

Different LLB's are designed for the Hi-Fi Sim and the Lo-Fi Sim. The LLB is capable of executing the currently selected HLB behaviour. In this particular embodiment, the Hi-Fi Sim LLB uses more computationally expensive
 30 techniques to execute the HLB behaviour. The Lo-Fi Sim LLB uses cheaper techniques that are designed to approximate the effect of the Hi-Fi Sim LLB as close as possible.

For example, the Hi-Fi Sim in this embodiment is arranged to use a form of three dimensional (3D) modelled environment, including 3D physical simulations of each
 5 animat. The Hi-Fi Sim LLB is consequently arranged to use neural networks to control the articulation of the limbs of the body of the animat, with the interaction of limbs with the environment being modelled using rigid-body dynamics for physical simulation.

10

On the other hand, the Lo-Fi Sim has only a 2D (two dimensional) model of the environment, with the corresponding LLB simply using a parameterised model of the Hi-Fi Sim locomotion to approximate the movement. For
 15 instance, the LLB treats the body of the animal as a point mass with predetermined parameters such as speed, rate of turn etc.

By separating the model of the animat into two portions,
 20 it is possible to ensure consistent animat behaviour in both Hi-Fi Sim and Lo-Fi Sim, while using the most appropriate techniques within each mode of simulation as required

25 Figure 1 shows the components of the simulation that are specific to the Hi-Fi Sim ("High Fidelity") and the Lo-Fi Sim ("Low Fidelity"). The dotted line 10 denotes the boundary between the simulation independent and the simulation specific components.

30

As can be seen, in the preferred embodiment the environment 20 is split into two portions, the LoFi environment abstraction 120 and the HiFi environment

implementation 220. There are correspondingly two representations of an animat: the LoFi Animat 30 and the HiFi Animat 40.

5 In both representations there are sensors 50 and actuators 60. In the LoFi representation the sensors are modelled by a computationally inexpensive abstraction 150 and the actuators are modelled by a similarly simplified abstraction 160. In the HiFi representation the sensors
10 are implemented using an accurate and realistic (but computationally expensive) implementation 250 and similarly the actuators are implemented using an accurate and realistic (but computationally expensive) implementation 260.

15

The sensors and actuators each require an appropriate representation of the simulated environment, and so in the LoFi Sim the process 1210 of computing what the animat senses in the environment 20 is achieved by use of an
20 abstract representation of the physics of sensing 121. Similarly the process 1220 of effecting the changes in the animat's relationship to its environment 20 as a consequence of its current actuator settings 160 is achieved by use of an appropriate abstract representation
25 of the physics of actuation 122.

In an analogous fashion for an instance of a HiFi Sim animat, the process 2210 of determining what the animat senses relies on the HiFi Sim sensors 250 having available
30 a high fidelity implementation of the relevant sensor-physics 221. Similarly the process 2220 of determining the effects of the HiFi Sim actuators 260 requires a high-

fidelity model of the physics of actuation 222 in the environment 20.

It should be understood that the word "physics" is used
5 here to refer to the nature of the environment's "laws of
physics" and can refer both to simulations of physics
systems similar to those familiar from common knowledge of
the laws of physics in the real known universe, and also
to fictional and imaginary physics systems unlike those
10 known to exist in the real universe. An example of an
imaginary physics system is the "laws of physics" of many
humorous animated cartoons, where an animated character
that runs at high speed off the edge of a cliff will
typically continue to run horizontally through the air
15 until it notices that it is in mid-air, at which point it
ceases to travel horizontally and instead starts to fall
vertically downwards.

Referring again to Figure 1, in an example of the
20 preferred embodiment of the invention the intended
application may be a computational simulation of wild
animals on the African plains. For each species of animal
that is simulated, there will be a LoFi Sim representation
and a corresponding HiFi Sim representation.

25 In this example animats are either simulated antelope or
simulated lions. Each individual animal in the simulation
will have its own LoFi Animat representation 30 that is
computationally inexpensive to calculate updates for and
30 that requires comparatively little storage space (RAM-
space or disk-space). A HiFi Animat representation of an
individual animal will only be created when necessary for

the purposes of the simulation, e.g. when the creature fights or pursues a prey or evades a predator.

In this example of the preferred embodiment, there may be
5 many thousands of simulated antelopes, and many hundreds
of lions, each of which is represented in the LoFi
Environment 120 as a point with co-ordinates (x,y) on a 2-
dimensional (2-d) plane. The LoFi Actuators 160 may simply
10 apply a 2-d vector to the animat's co-ordinates, updating
its position. The vector's direction and magnitude would
vary as the animat runs, eats, or sleeps. The antelope-
animat's LoFi sensors 150 may simply examine the LoFi
Sensor Physics representation of the environment 121 to
15 determine the distance and direction of the nearest food
source, the nearest water source, the nearest other
antelope and the nearest lion. Sources of food and water
may simply be represented in the LoFi Environment
Abstraction 120 as circular patches on the ground-plane,
and so can be represented by three numbers: the (x,y) co-
20 ordinates of the centre of the patch, and its radius. Thus
it can be seen that the LoFiSim representation of many
thousands of antelope, many hundreds of lions, and many
patches of food and many patches of water can be
represented without the need for much computation-power or
25 storage-space.

In this example it could be the case that if a lion is
sufficiently close to an antelope, if the lion is hungry
and if it has enough energy then it will chase the
30 antelope. At this point the LoFiSim abstraction 120 and
animat representations 30 are not sufficient to capture
the details of the interaction between the two creatures.
Consequently a HiFiSim representation of each animal is

created, with HiFi sensors 250 and HiFi actuators 260 interacting with the HiFi Environment Implementation 220.

For greater accuracy and realism, the HiFi Environment 220
5 may involve a 3-d model of the animals' environment, including representations of natural terrain features such as hills and trees and rivers that form potential obstacles or impediments. The HiFi Sensor Physics 221 may use 3D graphics techniques to model directional vision in
10 each HiFi animat, in which case the Sensor HiFi Implementaion would have to model the "retina" of each animat as an array of pixels with appropriate modelling of the animal's optics (field of view, acuity of vision, ect). Similarly the HiFi Actuator Physics 222 may use
15 advanced mathematical modelling of Newtonian kinematics and dynamics such as those known in the prior art "physics engine" commercial middleware software products "Mathengine" (found at mathengine.com) and "Havok" (found
20 at havok.com), in which case the Actuators HiFi Implementation 260 would represent the animal's body skeletal frame and musculature in a manner consistent with the physics engine chosen for the HiFi Actuator Physics 222.

25 It will be appreciated that the HiFi Actuator Physics 222 and the HiFi Sensor Physics 221 are more tightly coupled within the HiFi Environment Implementation 220 than are the corresponding aspects of the LoFi Environment Abstraction 120. Consider (for example) that what the
30 creature senses using vision is affected by the current orientation of the creature's eyes and head, which is in turn determined by the actions of the creature's gaze-actuators (eyeball and neck muscles). It is typically

necessary to represent such sensory-motor coupling in High Fidelity simulations but it can be abstracted away from in Low Fidelity representations.

5 It will furthermore be appreciated that the coupling of 221 and 222 within 220 can extend to the creature's "sensation of self", e.g. modelling an animal's interoception (e.g. sensation of pain/damage, satiation, etc) or the analogous monitoring functions in machines
10 (e.g. monitors/sensors of oil-levels, engine temperatures, tyre pressures, etc).

In the example being described here, the HiFi Animat instances of the chasing lion and the chased antelope will
15 clearly require comparatively large amounts of data-storage space and computer-processor calculations to resolve the outcome of the interaction (i.e. whether the antelope escapes or not). By only creating these HiFi Sim representations when they are needed, and destroying them
20 as soon as they are no longer needed, it is possible to give the impression of highly complex high-fidelity simulations of many thousands of creatures while using a computer processor and memory storage system that is actually capable of accurately computing and storing such
25 high-fidelity simulations for only a very small number of such simulated creatures. The HiFiSim Animat representations 40 are only created when necessary, and are only maintained for as long as is necessary. So long as only a few animats are interacting with their
30 environments (and the other animats) in ways requiring accurate and complex computations, the remainder of the creatures in the simulation can be satisfactorily represented using the LoFiSim Animat representation 30.

The LoFi Animat Representation 30 can be seen to include the High Level Brain or HLB 70 which is responsible for selecting particular behavioural states or modes of action (such as resting, sleeping, exploring, or eating). The transitions between these states are caused by changes in the creature's Internal State 90 (such as energy levels altering) and by changchanging the abstract sense-data provided by the LoFi Sensors Abstraction 150.

10

The HiFi Animat Representation 40 includes the HLB 70 and Internal State values 90 but additionally includes a Low Level Brain or LLB 80 which implements specific behaviours at the level of triggering actuators in a fashion appropriate for each behaviour. The actions of the LLB can be affected or modulated by the Internal State values 90 and by sensory information received from the Sensors HiFi Implementation 250. The actions of the HLB 70 can also affect the LLB 80. In each instance of the HiFi Animat Representation, the HLB may be affected by the Abstract LoFi sensory data coming from 150, but in this case the data source for 150 is the HiFi to LoFi converter 100.

Thus it can be appreciated that in this embodiment the Sensors LoFi Abstraction 150 and the HLB 70 are used in both the LoFi and HiFi complexity-levels of the simulation, while other components are specific solely to either the LoFi or the HiFi complexity-levels.

In the preferred embodiment, the HLB 70 is a simple stochastic action selection mechanism used in both the Hi-Fi Sim and the Lo-Fi Sim. The underlying mechanism is a Markov net, a finite-state process where the different

states represent different behavioural modes or activities. The probability of transitions from one state to another are dependent on the animat's internal state variables (i.e. simulated bio-chemical and energetic levels e.g. indicating the energy level or level of fear of the creature) and any sensory input it receives.

Figure 2 shows a state machine representing three possible states (rest 410, sleep 420 and explore 430) that are used in a simplistic behaviour selection mechanism of an artificial creature.

The arrows joining each of the states represent the probability of the transition between each state e.g. RS=0.3 indicates that the probability of the creature moving from the rest state to the sleep state is 0.3 for any given time interval, RE=0.5 indicates that the probability of the creature moving from the rest to the explore state is 0.5 in the same time interval, and RR=0.2 indicates that the probability of the creature remaining in the rest state is 0.2 in the same time interval.

At any one time, it is assumed that the creature is only ever in one of the possible states. Consequently, whenever the behaviour is reevaluated, only the probability of the behaviour changing from that state (either back to the same state, or to another valid state) is evaluated. The value of the probability of a transition from any given state (and either back to the same state or to another state) is between 0 and 1, with a value of 0 meaning there is no chance of transitioning between the specified states, and an output of 1 meaning that this is the only transition that is possible. Clearly, the sum of

all of the transition probabilities from a given state is always equal to 1.

In this preferred embodiment, the creature perceives its
5 environment via sensors. Preferably, such sensors can
identify information about the environment (which may be a
simulated real or imaginary physical environment). The
sensors detect, for example, the number of visible
10 objects, location, object type etc, with processing
occurring so as to determine whether the sensed items are
likely to affect the behaviour of the creature (has food
been sensed?, a rival?, a potential mate?).

Additionally, the creature has other inputs, corresponding
15 to internal drives, such as hunger, fatigue, thirst and
desire to reproduce.

The transition probability between each behavioural state
is not fixed but is instead a function dependent upon
20 factors such as sensory input, internal drives and the
temperament of the creature.

The transition probability values shown in figure 2 are
provided by way of example only, and would in fact vary
25 over time depending upon these other factors.

It will be appreciated that the above embodiment is
provided by way of example only, and that various other
embodiments will be apparent to the skilled person. For
30 instance, an embodiment of the invention might not be
applied to all of the creatures within a simulation, but
to only a portion of the creatures.

It will be appreciated that the simple state-transition network described here is only one possible embodiment of the HLB behaviour-selection mechanism, and other methods could be used instead. Examples of other methods include
 5 those described by:

- T. Tyrrell. *Computational Methods of Action Selection*. PhD Thesis, University of Edinburgh 1993.
- B. Blumberg. *Old Tricks, New Dogs: Ethology and Interactive Creatures*. PhD Thesis, Massachusetts
 10 Institute of Technology, 1997.

In addition to this, it can be appreciated that aspects of the approach described here in respect of simulation of
 15 creatures can be employed more generally to simulation. In particular, other situations may be considered in which a simulation may be split into a first part which defines states and the possible transitions between states - a state machine - and two or more alternative second parts
 20 which are adapted to implement these transitions at different levels of complexity.

A general "first part" is shown in Figure 3A, and closely resembles the state machine of Figure 2. Three states are
 25 shown (A, B and C), together with possible transitions between them, and (in this case), probabilities for each transition. The general "second part", or collection of second parts, is shown in Figure 3B. It is desirable first to have a rule evaluated 502 with reference to at
 30 least one condition determined 501 by the computational system. This may be, for example, processor time available to carry out the simulation or predicted to be available over a certain period. This may apply where the

simulation is a single simulation rather than one of a group, and that the constraint on the simulation is that it must be carried out within a fixed period of time despite having varying computational resources available.

5 The condition may also be an indication of the significance of the simulation (rather as in the case of animats - there may be some determination as to which simulated entities are important, or important at that time, and computational resources may be preferentially

10 allocated to such entities - note here, however, that the "simulated entities" may be simply processes rather than specifically creatures).

On the basis of the evaluation of the rule, alternative

15 second portion choices 503, 504, 505 are followed for evaluating each of the transitions in Figure 3A (as is shown here, there may be more than two alternative levels of complexity available, rather than simply two). Each of these second choices has the same set of transitions, but

20 different rules or methods for implementing these rules such that the choices are of different levels of computational complexity. Here, for example, choice 503 simply provides transition probabilities to be evaluated over a particular period of time. Choice 504 involves

25 further variables a, b and c in calculating transitions - this will be of greater computational complexity than for choice 503, but will, preferably, be such that choice 503 provides an approximation to the process of choice 504 (so that both can reasonably be regarded as implementations of

30 the first portion shown in Figure 3A). Choice 505 involves further variables d, e and f, and is thus presumed for this example only to be of still greater computational complexity - though it is appreciated that

the presence of additional variables will not necessarily result in computational complexity, and that this may also result from other causes (such as functions that are more complex to evaluate). These different choices are used
5 for simulation - Figure 3B shows a loop structure with periodic re-evaluation of the rule 502 after a waiting step 506 (suggesting that the simulation can change levels of complexity during the process - which may be desirable if, for example, the computational resources available
10 vary), but in other cases the simulation may simply continue until an end condition is reached.

As will be appreciated from the above, such an approach to simulation is not limited in its application to large-
15 scale simulation of multiple creatures. The approach can be applied to simulations of single creatures, entities or even processes with different levels of computational resource available, or may be applied to simulations of plural entities or processes.

20 The reader's attention is directed to all papers and documents which are filed concurrently with or previous to this specification in connection with this application and which are open to public inspection with this
25 specification, and the contents of all such papers and documents are incorporated herein by reference.

All of the features disclosed in this specification (including any accompanying claims, abstract and
30 drawings), and/or all of the steps of any method or process so disclosed, may be combined in any combination, except combinations where at least some of such features and/or steps are mutually exclusive.

Each feature disclosed in this specification (including any accompanying claims, abstract and drawings), may be replaced by alternative features serving the same, equivalent or similar purpose, unless expressly stated otherwise. Thus, unless expressly stated otherwise, each feature disclosed is one example only of a generic series of equivalent or similar features.

10 The invention is not restricted to the details of the foregoing embodiment(s). The invention extends to any novel one, or any novel combination, of the features disclosed in this specification (including any accompanying claims, abstract and drawings), or to any
15 novel one, or any novel combination, of the steps of any method or process so disclosed.